Performance Activities: Dynamic Load Balancing and Prophesy

Valerie Taylor, Xingfu Wu, Jonathan Geisler, Xin Li, Zhiling Lan

Northwestern University

Rick Stevens, Ivan Judson, Mark Hereld

Argonne National Laboratory

Michael Norman

University of California in San Diego

Greg Bryan

Oxford

Dynamic Load Balancing

- *ENZO: AMR*
- *Analyzed parallel performance*
  - *Dispersion*
  - *Dynamicity*
  - *Frequency of adaptations*
- *Developed a parallel DLB*
- *Extended DLB to distributed system*

Imbalance

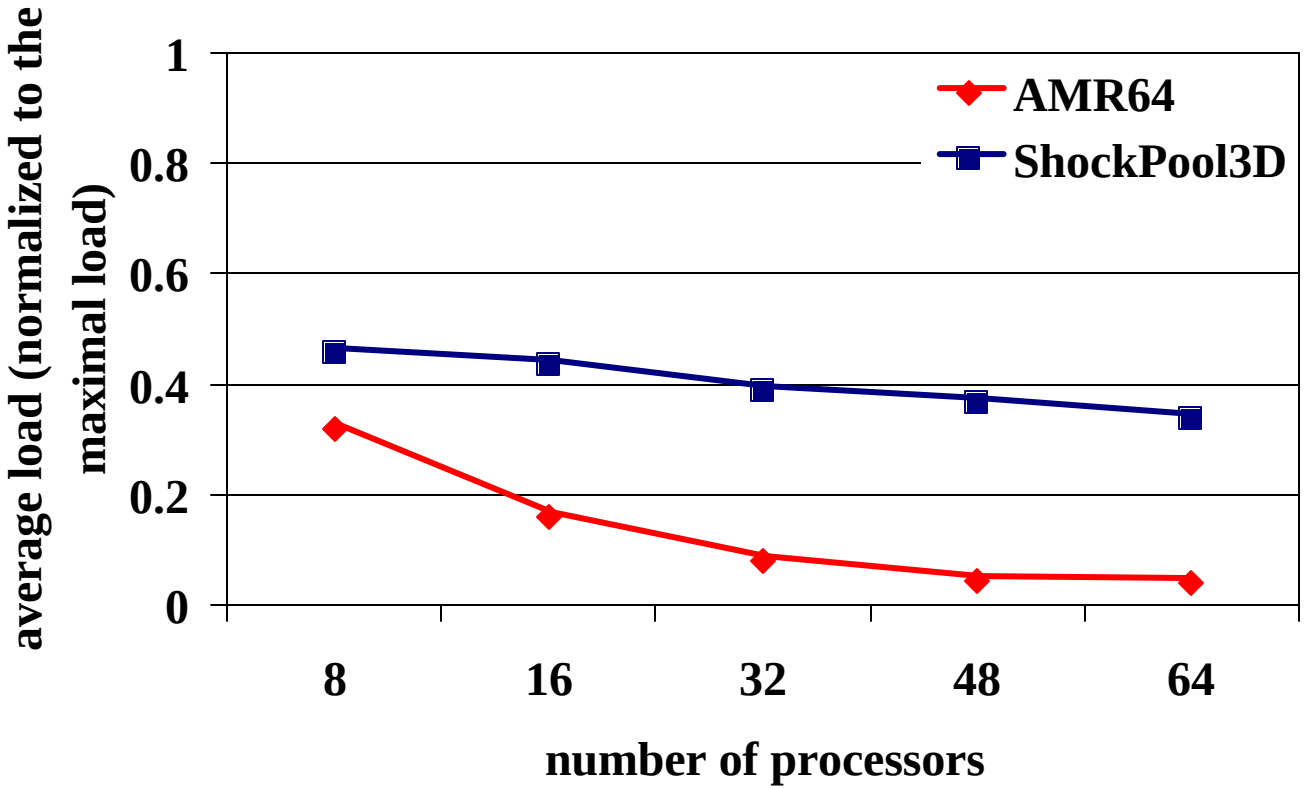

➡ ***The ideal case is close to 1.0. Imbalance exists for both datasets.***

Dispersion

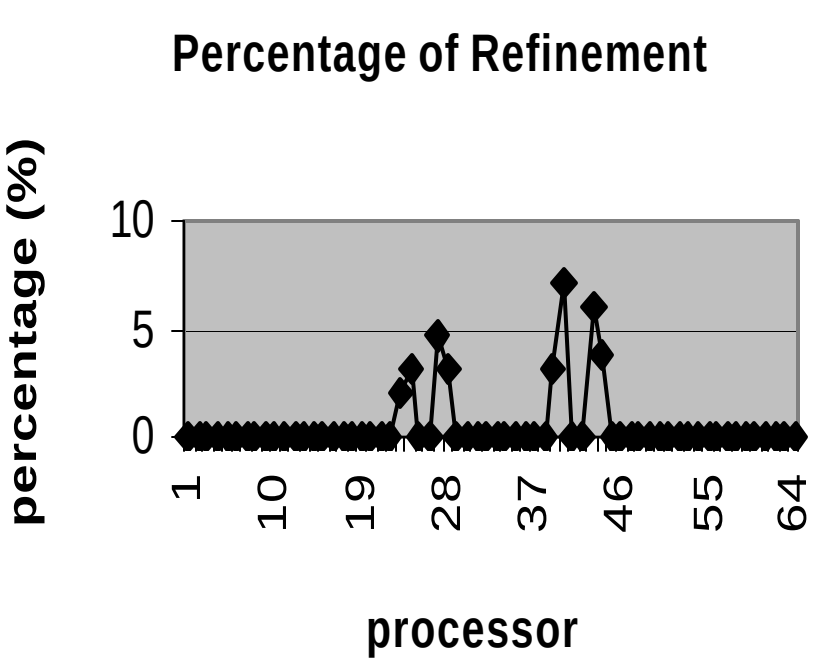

AMR64:
loads of a few processors are increased dramatically

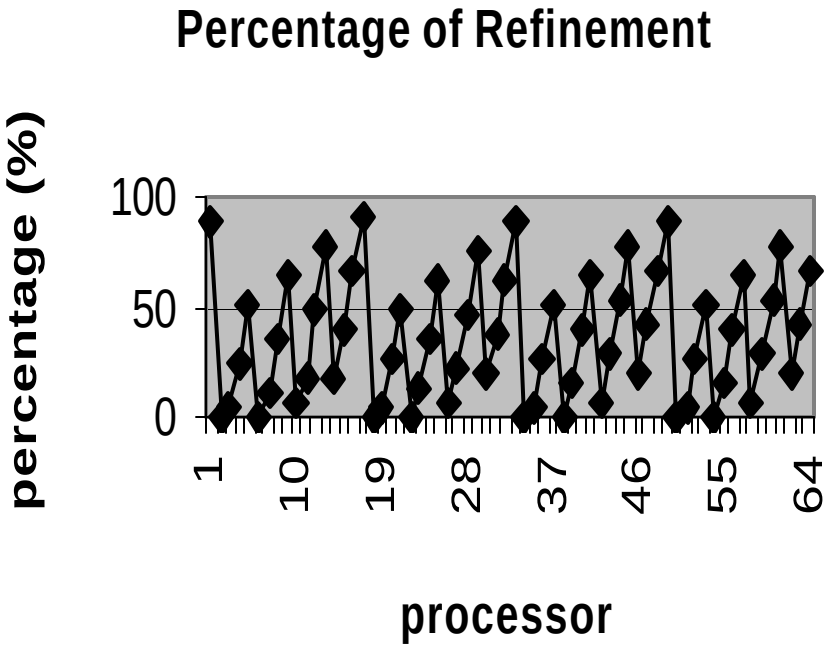

ShockPool3D:
there are four subgroups and each subgroup has the similar refinement characteristics

Execution Time for AMR64 on LAN-connect System

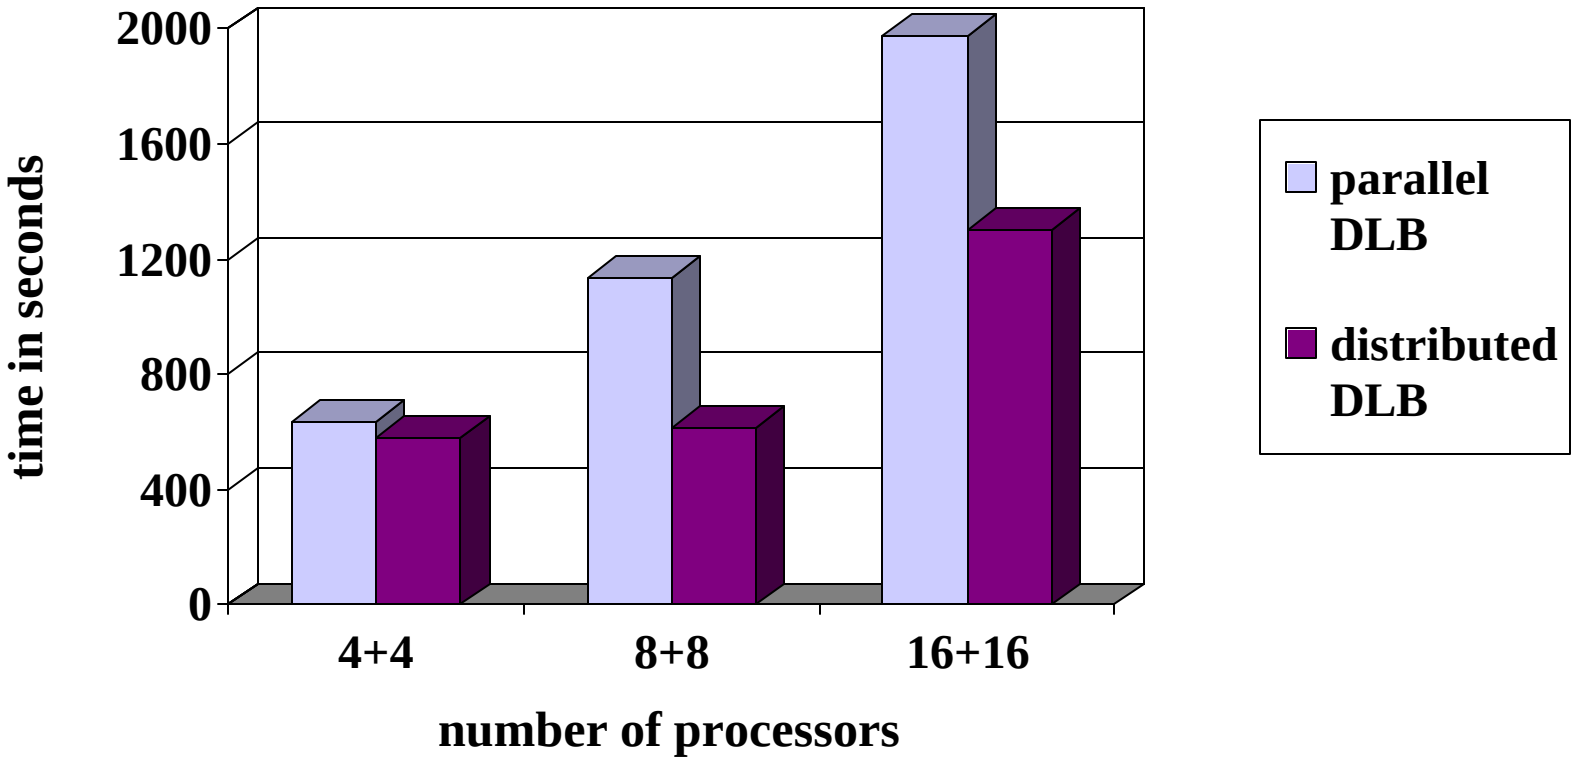

The relative improvement is ranging from 9.0% to 45.9%

Prophesy System

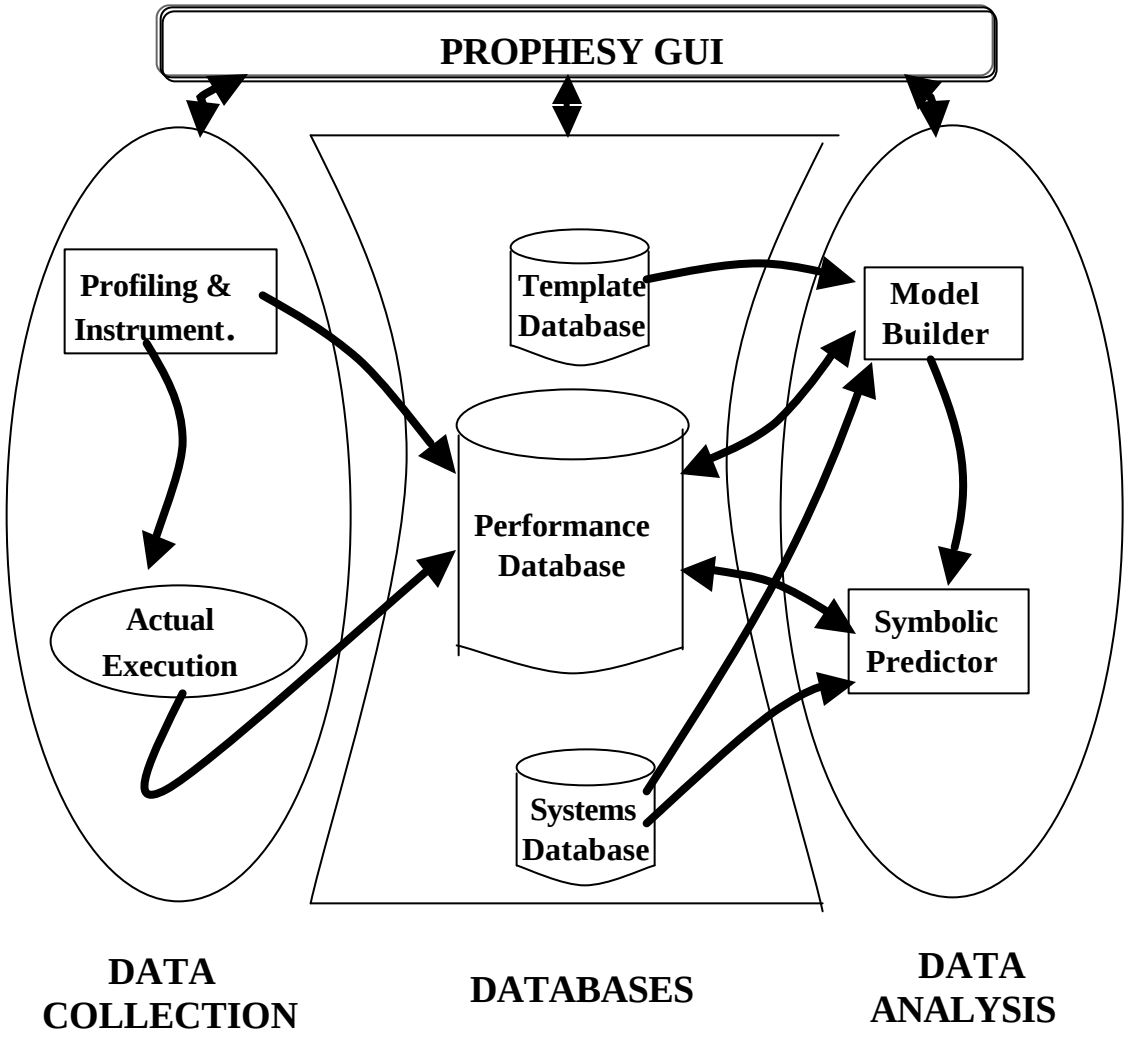

Automated Instrumentation

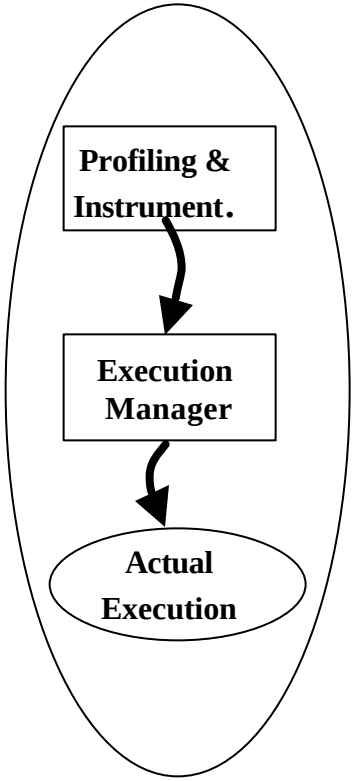

- *In-line data collection*
- *Instrument at the level of basic loops*
- *Allow for user-specified instrumentation*

```
T=E * f;
for (I=1; I<N; I++){
    V(I) = A(I) * C(I);
    B(I) = A(2I + 4);
}
```

```
T=E * f;
INSTRUMENTATION CODE
for (I=1; I<N; I++){
    V(I) = A(I) * C(I);
    B(I) = A(2I + 4);
}
INSTRUMENTATION CODE
```

Prophesy Database

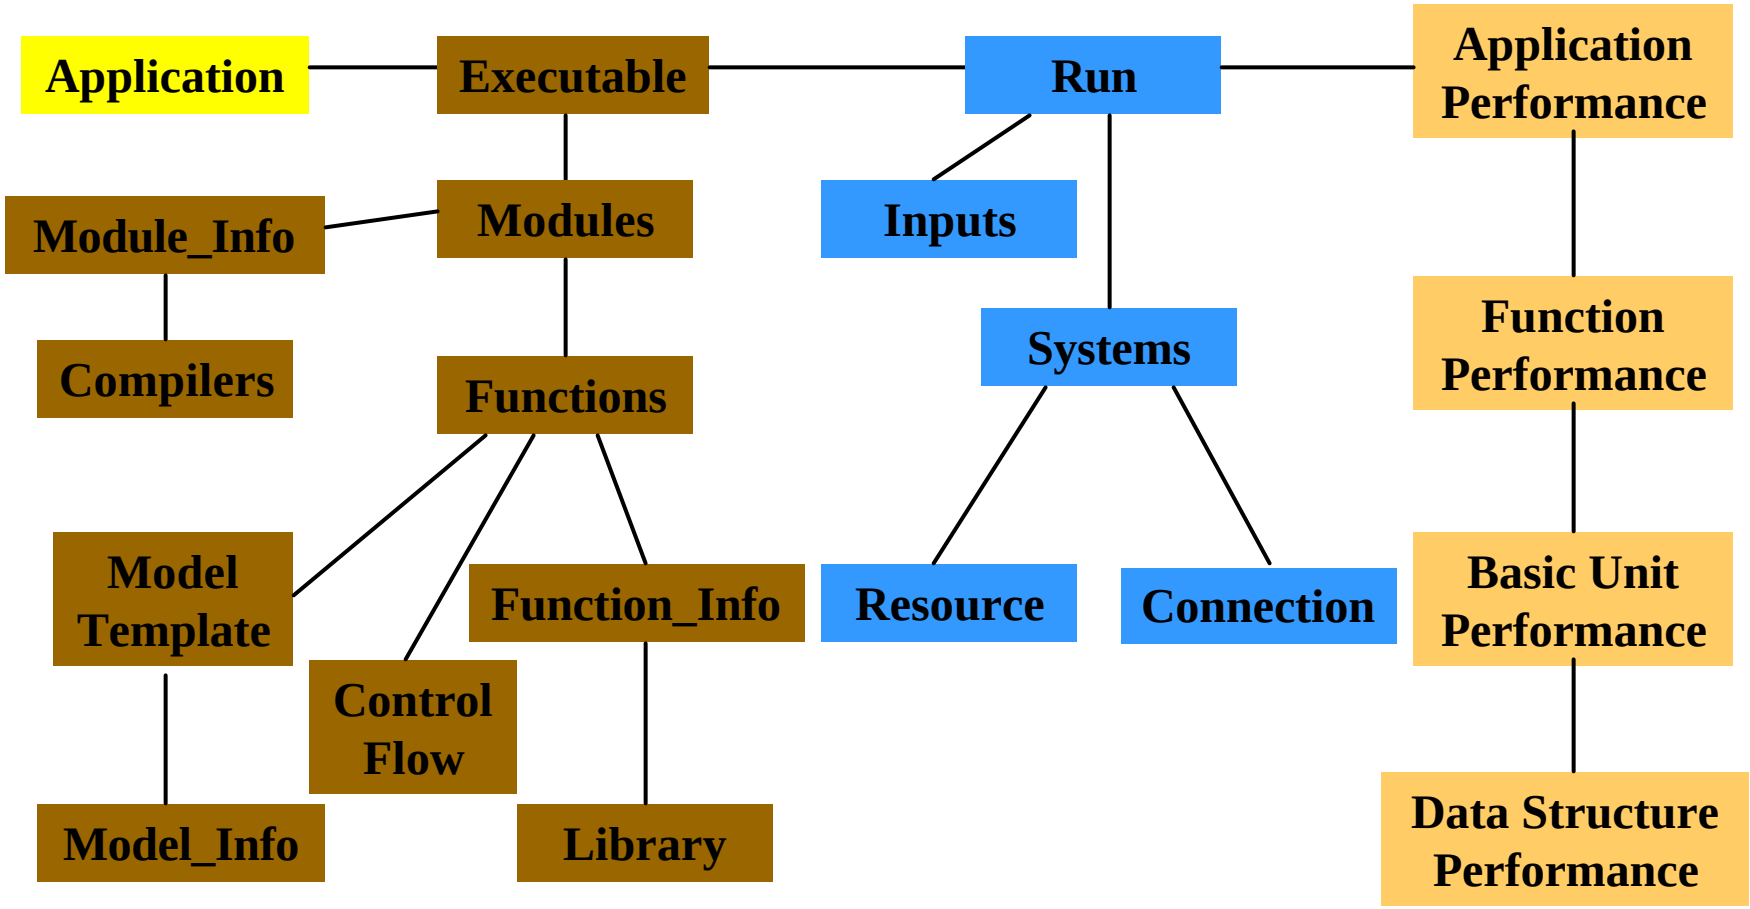

Data Analysis

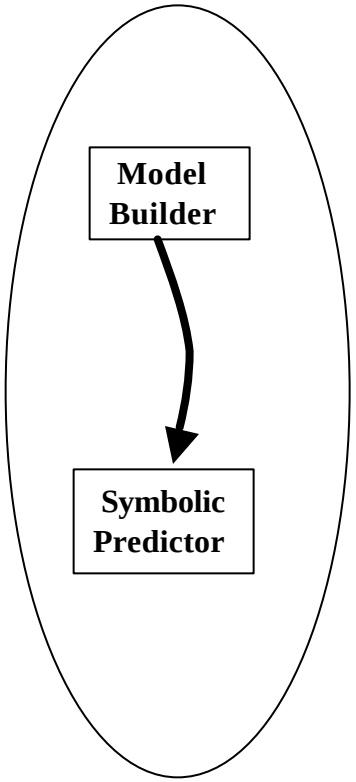

- Develop performance models
- Make predictions
- Performance tune codes
- Identify best implementation
- Identify trends

Look forward to collaborations

Automated Modeling Techniques

- *Utilize information in the template and system databases*
- *Currently include three techniques*
  - *Curve fitting*
  - *Parameterization*
  - *Composition using coupling values*

Parameterization: Example

- *Matrix-matrix multiply*
- *SGI (8 processors)*

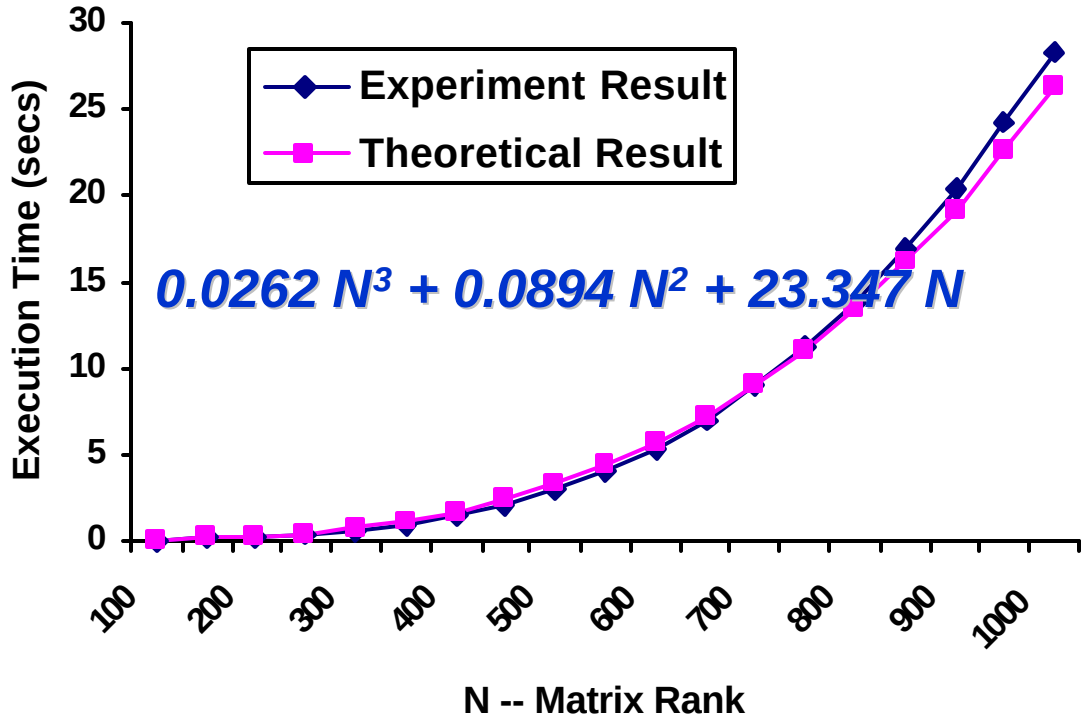

Dynamicity

- *Adaptation may be invoked after each time-step*
  - *for ShockPool3D, about 600 adaptations*
  - *for AMR64, more than 2500 adaptations*
  - *for AMR128, more than 5000 adaptations*
- *High frequency of adaptations*
  - *for AMR64, adaptation is invoked every 3 seconds*

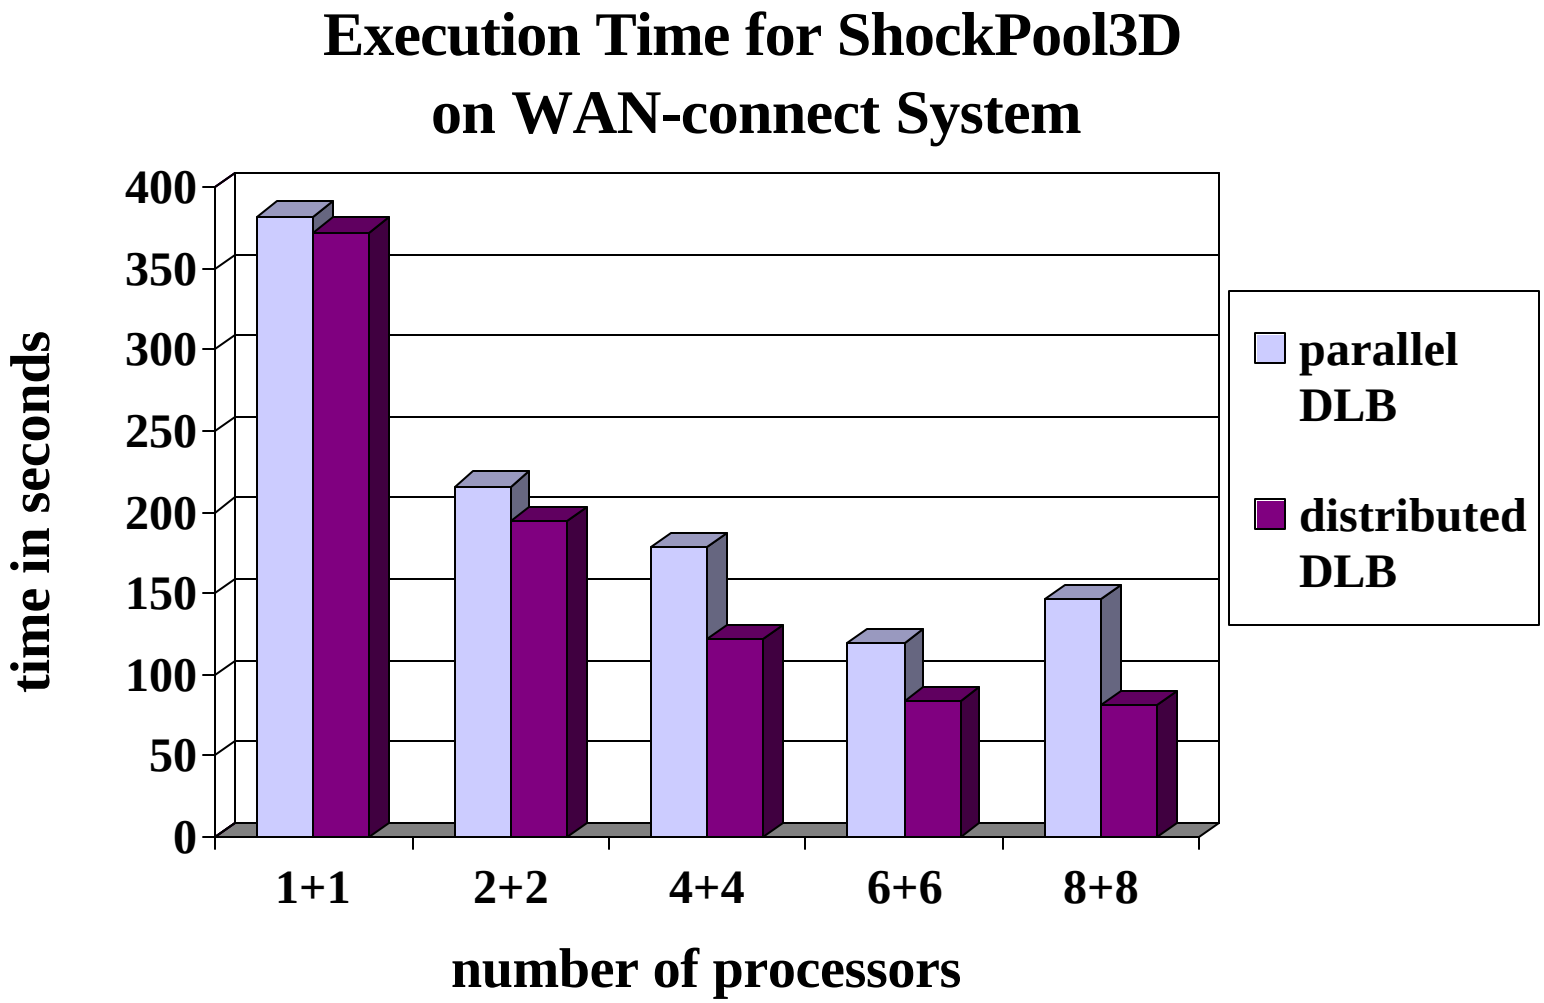

The relative improvement is ranging from 2.6% to 44.2%